TAPSpMV: Topology-Aware Parallel Sparse Matrix Vector Multiplication

Amanda Bienz, University of Illinois at Urbana-Champaign
William D. Gropp, University of Illinois at Urbana-Champaign
Luke N. Olson, University of Illinois at Urbana-Champaign

arXiv:1612.08060v1 [cs.DC] 23 Dec 2016

TODO

CCS Concepts: •**Theory of computation** → **Massively parallel algorithms;** *Parallel computing models;* •**Computing methodologies** → **Massively parallel algorithms;** •**Software and its engineering** → **Massively parallel systems;** •**Mathematics of computing** → *Solvers; Mathematical software performance;*

Additional Key Words and Phrases: Sparse, Matrix Vector Multiplication, SpMV, Parallel Communication, Topology Aware

ACM Reference Format:
Amanda Bienz, William D. Gropp, and Luke N. Olson, 2016. TAPSpMV: Topology-Aware Sparse Matrix Vector Multiplication *ACM Trans. Parallel Comput.* V, N, Article A (January YYYY), 18 pages.
DOI: http://dx.doi.org/10.1145/0000000.0000000

1. INTRODUCTION

The sparse matrix-vector multiply (SpMV) is a widely used operation in many simulations and the main kernel in iterative solvers. The focus of this paper is on the parallel SpMV, namely

$$w \leftarrow A * v \tag{1}$$

where A is a sparse $N \times N$ matrix and v is a dense N-dimensional vector. In parallel, the sparse system is often distributed across n_p processes such that each process holds a contiguous block of rows from the matrix A, and equivalent rows from the vectors v and w, as shown in Figure 1. A common approach is to also split the rows of A on a single process into two groups: an on-process block, containing the columns of the matrix that correspond to vector values stored locally, and an off-process block, containing matrix non-zeros that are associated with vector values that are stored on non-local processes. Therefore, non-zeros in the off-process block of the matrix require vector values to be communicated during each SpMV.

The SpMV operation lacks parallel scalability due to large costs associated with communication, specifically in the strong scaling limit of a few rows per process. In-

This research is part of the Blue Waters sustained-petascale computing project, which is supported by the National Science Foundation (awards OCI-0725070 and ACI-1238993) and the state of Illinois. Blue Waters is a joint effort of the University of Illinois at Urbana-Champaign and its National Center for Supercomputing Applications.
Author's addresses: A. Bienz, Department of Computer Science, University of Illinois at Urbana-Champaign; W. D. Gropp, Department of Computer Science, University of Illinois at Urbana-Champaign, http://wgropp.cs.illinois.edu; L. N. Olson, Department of Computer Science, University of Illinois at Urbana-Champaign, http://lukeo.cs.illinois.edu.
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
© YYYY ACM. 1539-9087/YYYY/01-ARTA $15.00
DOI: http://dx.doi.org/10.1145/0000000.0000000

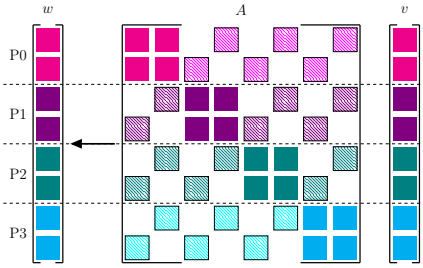

Fig. 1: A matrix partitioned across four processes, where each process stores two rows of the matrix, and the equivalent rows of each vector. The on-process block of each matrix partition is represented by solid squares, while the off-process block is represented by patterned entries.

creasing the number of processes for a matrix increases the number of columns in the off-process blocks, yielding a growth in communication.

Figure 2 shows the percentage of time spent communicating during a SpMV operation for two large matrices from the University of Florida sparse matrix collection[1] at scales varying from 50 000 to 500 000 non-zeros per process. The results show that the communication time becomes dominates the computation as the number of processes is increased, thus decreasing the scalability.

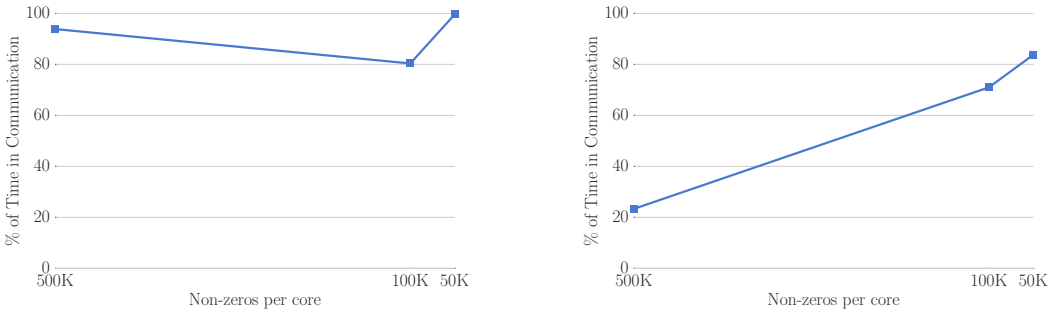

(a) `hv15r`: 283,073,458 non-zeros (b) `nlpkkt240`: 760,648,352 non-zeros

Fig. 2: Percentage of total SpMV time spent during communication.

Machine topology plays an important role in the cost of communication. Multicore distributed systems present new challenges in communication as the bandwidth is limited while the number of cores participating in communication increases [Gropp et al. 2016]. Injection limits and network contention are significant roadblocks in the SpMV operation, motivating the need for SpMV algorithms that take advantage of the machine topology. The focus of the approach developed in this paper is to use the node-processor hierarchy to more efficiently map communication, leading to notable reductions in SpMV costs on modern HPC system for a range of sparse matrix patterns.

There are a number of existing approaches for reducing communication costs associated with sparse matrix-vector multiplication. Communication volume in particular

[1] See http://www.cise.ufl.edu/research/sparse/matrices

is a limiting factor and the ordering and parallel partition of a matrix both influence the total data volume. In response, graph partitioning techniques are used to identify more efficient layouts in the data. ParMETIS [Karypis and Kumar 1998] and PT-Scotch [Chevalier and Pellegrini 2008], for example, provide parallel partitioning of matrices that often lead to improved system loads and more efficient sparse matrix operations. Communication volume is accurately modeled through the use of a hypergraph. As a result, hypergraph partitioning also leads to a reduction in parallel communication requirements, albeit at a larger one-time setup cost. The approach introduced in this paper complements these efforts by providing an additional level of optimization in handling communication.

This paper focuses on increasing the locality of communication during a SpMV to reduce the amount of communication injected into the network. Section 2 describes a reference algorithm for a parallel SpMV, which resembles the approach commonly used in practice. A performance model is also introduced, which considers the cost of intra- and inter-node communication and the impact on performance. A new SpMV algorithm is presented in Section 4, which reduces the number and size of inter-node messages by increasing the significantly cheaper intra-node communication. The code and numerics are presented in Section 5 to verify the performance.

2. BACKGROUND

Modern supercomputers incorporate a large number of nodes through an interconnect to form a multi-dimensional grid or torus network. Standard compute nodes are comprised of one or more multicore processors that share a large memory bank. The algorithm developed in this paper targets a general machine with this layout and the results are highlighted on two such systems, each with a different interconnect: Blue Waters, a Cray machine at the National Center for Supercomputing Applications, and Mira, an IBM Blue Gene Q at Argonne National Laboratory. Blue Waters consists of $22\,640$ Cray XE nodes, each containing two AMD 6276 Interlagos processors for a total of 16 cores per node[2]. The XE nodes are connected through a three-dimensional torus Gemini interconnect, with each Gemini consisting of two nodes. In contrast, Mira consists of $49\,152$ nodes each with a 16-core PowerPC processor connected through a five-dimensional torus[3].

To start, consider a system with n_p processes distributed across n_n nodes, resulting in $\texttt{ppn}$ processes per node. Rank $r \in [0, n_p - 1]$ is described by the tuple (p, n) where $0 \leq p < \texttt{ppn}$ is the local process number of rank r on node n. Assuming SMP-style ordering, the first $\texttt{ppn}$ ranks are mapped to the first node, the next $\texttt{ppn}$ to the second node, and so on. Therefore, rank r is described by the tuple $\left(r \mod \texttt{ppn}, \lfloor \frac{r}{\texttt{ppn}} \rfloor\right)$. Thus, for the remainder of the paper, the notation of rank r is interchangeable with $n \cdot \texttt{ppn} + p$.

Parallel matrices and vectors are distributed across all n_p ranks such that each process holds a portion of the linear system. Let $\mathcal{R}(r)$ be the rows of an $N \times N$ sparse linear system, $w \leftarrow A * v$, stored on rank r. In the case of an even, contiguous partition where the k^{th} partition is placed on the k^{th} rank, $\mathcal{R}(r)$ is defined as

$$\mathcal{R}(r) = \left\{ \left\lfloor \frac{N}{n_p} \right\rfloor r, \ldots, \left\lfloor \frac{N}{n_p} \right\rfloor (r+1) - 1 \right\} \tag{2}$$

[2]https://bluewaters.ncsa.illinois.edu/hardware-summary
[3]http://www.alcf.anl.gov/user-guides/mira-cetus-vesta

or equivalently as

$$\mathcal{R}(n \cdot \texttt{ppn} + p) = \left\{ \left\lfloor \frac{N}{n_p} \right\rfloor (n \cdot \texttt{ppn} + p), \ldots, \left\lfloor \frac{N}{n_p} \right\rfloor (n \cdot \texttt{ppn} + p + 1) - 1 \right\}. \quad (3)$$

An example parallel system is displayed in Figure 3, consisting of six processes distributed across three nodes. The matrix in Figure 4 is partitioned across this processor layout with each process holding a single row of the matrix and associated row of the input vector. In this example, the diagonal entry falls into the on-process block, as the corresponding vector value is stored locally. The off-process block, which requires communication, consists of all off-diagonal non-zeros as the associated vector values are stored on other processes. The off-process block of A is further partitioned into on-node

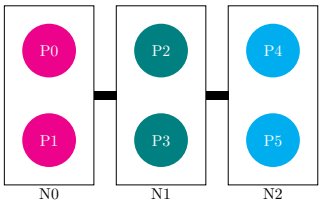

Fig. 3: An example parallel system with six processes distributed across three nodes.

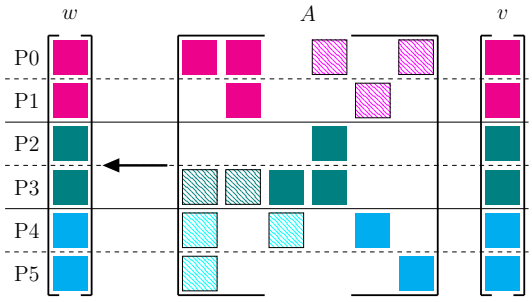

Fig. 4: An example 6×6 sparse matrix for the parallel system in Figure 3. The solid shading denotes blocks that require only on-node communication, while the striped shading denotes blocks that require communication with distant nodes.

and off-node blocks, as labeled in Figure 4. For any process (p, n), the on-node columns of A correspond to vector values that are stored on some process (s, n), where $s \neq p$. Similarly, the off-node columns of A, correspond to data stored on some process (q, m), where $m \neq n$. To make this clearer, we define the following

$$\texttt{on_process}(A, n \cdot \texttt{ppn} + p) = \{a_{ij} \neq 0 \,|\, i, j \in R(n \cdot \texttt{ppn} + p)\} \quad (4)$$

$$\texttt{off_process}(A, n \cdot \texttt{ppn} + p) = \{a_{ij} \neq 0 \,|\, i \in R(n \cdot \texttt{ppn} + p), j \notin R(n \cdot \texttt{ppn} + p)\} \quad (5)$$

$$\texttt{on_node}(A, n \cdot \texttt{ppn} + p) = \{a_{ij} \neq 0 \,|\, \exists\, q \neq p \text{ with } i \in R(n \cdot \texttt{ppn} + p), j \in R(n \cdot \texttt{ppn} + q)\} \quad (6)$$

and

$$\texttt{off_node}(A, n \cdot \texttt{ppn} + p) = \{a_{ij} \neq 0 \,|\, \exists\, q, m \neq n \text{ with } i \in R(n \cdot \texttt{ppn} + p),\, j \in R(m \cdot \texttt{ppn} + q)\}\,. \quad (7)$$

2.1. Standard SpMV

For a sparse matrix-vector multiply, $w \leftarrow A * v$, each process receives all values of v associated with the non-zero entries in the off-process block of A. For example, if rank r contains a non-zero entry of A, A_{ij}, at row i, column j, then rank s with row $j \in \mathcal{R}(s)$ sends the j^{th} vector value, v_j, to rank r. Typically, these communication requirements are determined as the sparse matrix is formed.

In the standard SpMV, for each rank r there are lists of processes to which data is sent and from which data is received, as well as the global vector indices for each of these processes. The functions $\mathcal{P}_{\text{send}}(r)$ and $\mathcal{P}^{\text{recv}}(s)$ define these lists, and P_s^r indicates whether a message is sent from process r to s. Specifically,

$$P_s^r = 1 \quad \text{if } \exists\, A_{ij} \neq 0 \text{ with } i \in \mathcal{R}(r),\, j \in \mathcal{R}(s) \quad (8)$$

$$\mathcal{P}_{\text{send}}(r) = \{s \,|\, P_s^r = 1\}\,, \quad (9)$$

$$\mathcal{P}^{\text{recv}}(s) = \{r \,|\, P_s^r = 1\}\,. \quad (10)$$

As an example, consider a SpMV for the linear system in Figure 4, which is partitioned across an architecture described in Figure 3. Each column in Table I indicate the processes to which each rank r sends messages, while the rows indicate the processes from which each rank s receives messages. The values in $\mathcal{P}_{\text{send}}(r)$ and $\mathcal{P}^{\text{recv}}(r)$ are further displayed in Table II.

		r					
		0	1	2	3	4	5
s	0	0	1	0	1	0	1
	1	0	0	0	0	1	0
	2	0	0	0	1	0	0
	3	1	1	1	0	0	0
	4	1	0	1	0	0	0
	5	1	0	0	0	0	0

Table I: Communication pattern for ranks r and s for the example given in Figure 4. The value of P_s^r is given for each combinations of ranks.

	r					
	0	1	2	3	4	5
$\mathcal{P}^{\text{send}}(r)$	$\{3,4,5\}$	$\{0,3\}$	$\{3,4\}$	$\{0,2\}$	$\{1\}$	$\{0\}$
$\mathcal{P}_{\text{recv}}(r)$	$\{1,3,5\}$	$\{4\}$	$\{3\}$	$\{0,1,2\}$	$\{0,2\}$	$\{0\}$

Table II: Communication pattern for rank r for the example given in Figure 4, containing the values for $\mathcal{P}_{\text{send}}(r)$ and $\mathcal{P}^{\text{recv}}(r)$ for each rank r.

.

For each s in $\mathcal{P}_{\text{send}}(r)$, define the function $\mathcal{D}(r, s)$ to return the global vector indices that process r sends to process s. In addition, this function $\mathcal{D}(r, s)$ describes the vector

indices that s receives from r, for each r in $\mathcal{P}^{\text{recv}}(s)$. This function is defined as the following.

$$\mathcal{D}(r,s) = \{i \,|\, A_{ij} \neq\ 0 \text{ with } i \in \mathcal{R}(r),\, j \in \mathcal{R}(s)\} \tag{11}$$

Following the example in Figure 3, Table III lists the indices that each rank r sends to each rank s. Similarly, the rows list the indices s that receive from each rank r some rank s receives.

		r					
		0	1	2	3	4	5
s	0	{}	{1}	{}	{3}	{}	{5}
	1	{}	{}	{}	{}	{4}	{}
	2	{}	{}	{}	{3}	{}	{}
	3	{0}	{1}	{2}	{}	{}	{}
	4	{0}	{}	{2}	{}	{}	{}
	5	{0}	{}	{}	{}	{}	{}

Table III: Each column r lists the values sent to each process s in $\mathcal{P}_{\text{send}}(r)$, namely $\mathcal{D}(r,s)$. Each row s lists the values received from each process r in $\mathcal{P}^{\text{recv}}(s)$, namely $\mathcal{D}(r,s)$

.

With these definitions, the *standard* or reference SpMV is described in Algorithm 1. It is important to note that the parallel communication in Algorithm 1 is executed independent of any locality in the problem. That is, messages sent to another process may be both on-node or off-node depending on the process, however this is not considered in the algorithm.

ALGORITHM 1: `standard_spmv`

Input: r
$A|_{\mathcal{R}(r)}$
$v|_{\mathcal{R}(r)}$

Output: $w|_{\mathcal{R}(r)}$

$A_{\text{on_process}} = \texttt{on_process}(A|_{\mathcal{R}(r)})$
$A_{\text{off_process}} = \texttt{off_process}(A|_{\mathcal{R}(r)})$
for $s \in \mathcal{P}_{send}(r)$ **do**
  for $i \in \mathcal{D}(r,s)$ **do**
    $b_{\text{send}} \leftarrow v|_{\mathcal{R}(r)_i}$
  end
  `MPI_Isend`$(b_{\text{send}}, \ldots, s, \ldots)$
end
$b_{\text{recv}} \leftarrow \emptyset$
for $s \in \mathcal{P}_{recv}(r)$ **do**
  `MPI_Irecv`$(b_{\text{recv}}, \ldots, s, \ldots)$
end
`sequential_spmv`$(A_{\text{on_process}}, v|_{\mathcal{R}(r)})$
`MPI_Waitall`
`sequential_spmv`$(A_{\text{off_process}}, b_{\text{recv}})$

3. COMMUNICATION MODELS

The performance of Algorithm 1 is sub-optimal since it does not take advantage of locality in the communication. To see this, a communication performance model is

	α	B_{inj}	$B_{\max}$	B_{N}
Short	$4.0 \cdot 10^{-6}$	$6.3 \cdot 10^{8}$	$-1.8 \cdot 10^{7}$	∞
Eager	$1.1 \cdot 10^{-5}$	$1.7 \cdot 10^{9}$	$6.2 \cdot 10^{7}$	∞
Rendezvous	$2.0 \cdot 10^{-5}$	$3.6 \cdot 10^{9}$	$6.1 \cdot 10^{8}$	$5.5 \cdot 10^{9}$

Table IV: Measurements for α, B_{inj}, $B_{\min}$, and B_{N} for Blue Waters.

developed in this section. One approach is that of the *max-rate model*[Gropp et al. 2016], which describes the time-to-communicate as

$$T = \alpha + \frac{\texttt{ppn} \cdot s}{\min(B_{\text{N}}, B_{\max} + (\texttt{ppn} - 1)B_{\text{inj}})} \tag{12}$$

where α is the *latency* or start-up cost of a message, which may include preparing a message for transport or determining the network route; s is the number of bytes to be communicated; ppn is again the number of communicating processes per node; B_{inj} is the maximum rate at which messages are injected into the network; $B_{\max}$ is the achievable message rate of each process or *bandwidth*; and B_{N} is the peak rate of the NIC. In the simplest case of $\texttt{ppn} = 1$, the familiar *postal model* suffices:

$$T = \alpha + \frac{s}{B_{\max}}. \tag{13}$$

MPI contains multiple message passing protocols, including short, eager, and rendezvous. Each message consists of an envelope, including information about the message such as message size and source information, as well as message data. Short messages contain very little data which is sent as part of the envelope. Eager and rendezvous messages, however, send the envelope followed by packets of data. Eager messages are sent under the assumption that the receiving process has buffer space available to store data that is communicated. Therefore, a message is sent without checking buffer space at the receive process, limiting the associated latency. However, if a message is sufficiently large, rendezvous protocol must be used. This protocol requires the sending process to inform the receiving rank of the message so that buffer space is allocated. The message is sent only once the sending process is informed that this space is available. Therefore, there is a larger overhead with sending a message using rendezvous protocol. Table IV displays the measurements for α, B_{inj}, $B_{\max}$, and B_{N} for Blue Waters.

The *max-rate* model can be improved by distinguishing between intra- and inter-node communication. If the sending and receiving processes lie on the same physical node, data is not injected into the network, yielding low start-up and byte transport costs. As intra-node messages are not injected into the network, communication local to a node can be modeled as

$$T_\ell = \alpha_\ell + \frac{s_\ell}{B_{\max_\ell}} \tag{14}$$

where α_ℓ is the start-up cost for intra-node messages; s_ℓ is the number of bytes to be transported; and $B_{\max_\ell}$ is the achievable intra-node message rate.

Nodecomm[4] [Gropp et al. 2016], a topology-aware communication program, measures the time required to communicate on various levels of the parallel system, such as between two nodes of varying distances and between processes local to a node. Communication tests between processes local to one node were used ot caluclate the intra-node model parameters, as displayed in Table V.

[4]See https://bitbucket.org/william_gropp/baseenv

	α_ℓ	$B_{\max_\ell}$
Short	$1.3 \cdot 10^{-6}$	$4.2 \cdot 10^{8}$
Eager	$1.6 \cdot 10^{-6}$	$7.4 \cdot 10^{8}$
Rendezvous	$4.2 \cdot 10^{-6}$	$3.1 \cdot 10^{9}$

Table V: Measurements for intra-node variables, α_ℓ and $B_{\max_\ell}$.

Furthermore, Figure 5 shows the the time required to send a single message of varying sizes. The thin lines display nodecomm measurements for time required to send a single message, as either inter- or intra-node communication. Furthermore, the thick lines represent the time required to send a message of each size, according to the *max-rate* model in (12) and intra-node model in (14). This figure displays a significant difference between the costs of intra- and inter-node communication.

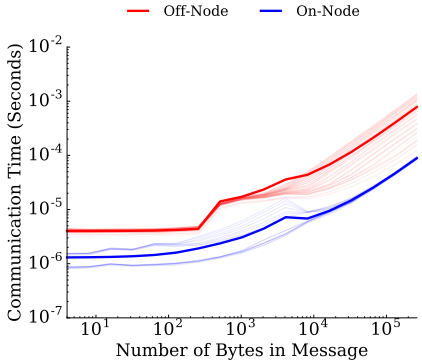

Fig. 5: The time required to send a single message of various sizes, with the thin lines representing timings measured by nodecomm and the thick lines displaying the *max-rate* and intra-node models in (12) and (14), respectively.

4. TAPSPMV

To reduce communication costs, the algorithm proposed in this section decreases the number and size of messages being injected into the network by increasing the amount of intra-node communication, which is less-costly than inter-node communication. This trade-off is accomplished through a so-called *topology-aware parallel* SpMV (TAPSpMV), where values are gathered in processes local to each node before being sent across the network, followed by a distribution of processes on the receiving node. As a result, each process (p, n) determines the communicating processes during the various steps of a TAPSpMV, as well as the accompanying data.

4.1. Inter-node communication setup

To eliminate the communication of duplicated messages, a list of communicating nodes is formed for each node n along with the accompanying data values. These lists are then distributed across all processes local to n by balancing the number of nodes and volume of data for communication. To facilitate this, define $N_m^n = 1$ if any process p on node n communicates with a process q on node m. In addition, the functions $\mathcal{N}_{\text{send}}(n)$

and $\mathcal{N}^{\text{recv}}(n)$ define the set of communication nodes for processes on node n. Specifically,

$$N_m^n = 1 \text{ if } \exists\, p, q \text{ s.t. } A_{ij} \neq 0 \text{ with } i \in \mathcal{R}(n \cdot \texttt{ppn} + p),\, j \in \mathcal{R}(m \cdot \texttt{ppn} + q),\, n \neq m \tag{15}$$

$$N_n^n = 0 \,\forall\, n \tag{16}$$

$$\mathcal{N}_{\text{send}}(n) = \{m \,|\, N_m^n = 1\} \tag{17}$$

$$\mathcal{N}^{\text{recv}}(m) = \{n \,|\, N_m^n = 1\} \tag{18}$$

Table VI displays N_m^n for the example in Figure 3. In addition, Table VII contains the lists of communicating nodes for node n: $\mathcal{N}_{\text{send}}(n)$ and $\mathcal{N}^{\text{recv}}(n)$.

		n = 0	1	2
m	0	0	1	1
	1	1	0	0
	2	1	1	0

Table VI: Communication requirements for nodes n and m for the example in Figure 4. Each column n contains the values in $\mathcal{N}_m^n$; each each row m contains the values in $\mathcal{N}_m^n$.

.

	n = 0	1	2
$\mathcal{N}^{\text{send}}(n)$	$\{1,2\}$	$\{0,2\}$	$\{0\}$
$\mathcal{N}_{\text{recv}}(n)$	$\{1,2\}$	$\{0\}$	$\{0,1\}$

Table VII: Communication pattern for each node n for the example in Figure 4.

The associated data values are defined for each node $m \in \mathcal{N}_{\text{send}}(n)$ with $\mathcal{E}(n,m)$, which returns the data indices to be sent from node n to node m. In other words, for each node $n \in \mathcal{N}^{\text{recv}}(m)$, $\mathcal{E}(n,m)$ defines the set of data indices that m receives from n. That is,

$$\mathcal{E}(n,m) = \{i \,|\, \exists\, p, q \text{ s.t. } A_{ij} \neq 0 \text{ with } i \in \mathcal{R}(n \cdot \texttt{ppn} + p),\, j \in \mathcal{R}(m \cdot \texttt{ppn} + q),\, n \neq m\}\,. \tag{19}$$

Extending the example in Figure 3, Table VIII displays the global vector indices, $\mathcal{E}(n,m)$, for each set of nodes n and m.

		n = 0	1	2
m	0	$\{\}$	$\{3\}$	$\{4,5\}$
	1	$\{0,1\}$	$\{\}$	$\{\}$
	2	$\{0\}$	$\{2\}$	$\{\}$

Table VIII: For the example given in Figure 4, each column n contains the values sent from n to m, while each row m contains values that node m receives from node n, as in $\mathcal{E}(n,m)$.

$\mathcal{T}_{\text{send}}(n \cdot \texttt{ppn} + p)$ defines the nodes that are distributed to process (p, n), which describes he mapping to distribute $\mathcal{N}_{\text{send}}(n)$ across all processes p on node n. With a similar definition for the *receiving* nodes, these are

$$\mathcal{T}_{\text{send}}(n \cdot \texttt{ppn} + p) = \{m \in \mathcal{N}_{\text{send}}(n) \,|\, m \text{ maps to } (p, n)\}, \tag{20}$$
$$\mathcal{T}^{\text{recv}}(m \cdot \texttt{ppn} + q) = \{n \in \mathcal{N}^{\text{recv}}(m) \,|\, n \text{ maps to } (q, m)\}. \tag{21}$$

The processor layout in the example in Figure 3 is displayed in Table IX, where the columns contain the send and receive nodes that are mapped to each process.

	(p, n)					
	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
$\mathcal{T}_{\text{send}}(n \cdot \texttt{ppn} + p)$	$\{1\}$	$\{2\}$	$\{0\}$	$\{2\}$	$\{0\}$	$\{\}$
$\mathcal{T}^{\text{recv}}(n \cdot \texttt{ppn} + p)$	$\{2\}$	$\{1\}$	$\{\}$	$\{0\}$	$\{1\}$	$\{0\}$

Table IX: Processor mappings for $\mathcal{N}_{\text{send}}(n)$ and $\mathcal{N}^{\text{recv}}(n)$ for the example given in Figure 4.

Finally, $\mathcal{G}_{\text{send}}(n \cdot \texttt{ppn} + p)$ defines the set of all off-node processes that process (p, n) sends to during the inter-node communication step of the topology-aware SpMV. Specifically, with a similar definition for receiving off-process data,

$$\mathcal{G}_{\text{send}}(n \cdot \texttt{ppn} + p) = \{(p, n) \,|\, m \in \mathcal{T}_{\text{send}}(n \cdot \texttt{ppn} + p),\, n \in \mathcal{T}^{\text{recv}}(m \cdot \texttt{ppn} + q)\}, \tag{22}$$
$$\mathcal{G}^{\text{recv}}(m \cdot \texttt{ppn} + q) = \{(q, m) \,|\, n \in \mathcal{T}^{\text{recv}}(m \cdot \texttt{ppn} + q),\, m \in \mathcal{T}_{\text{send}}(n \cdot \texttt{ppn} + p)\}. \tag{23}$$

Following the example in Figure 3, Table X displays the processes to which each (p, n) sends, $\mathcal{G}_{\text{send}}(n \cdot \texttt{ppn} + p)$, and from which each (p, n) receives, $\mathcal{G}^{\text{recv}}(n \cdot \texttt{ppn} + p)$. Finally,

	(p, n)					
	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
$\mathcal{G}_{\text{send}}(n \cdot \texttt{ppn} + p)$	$\{(1,1)\}$	$\{(1,2)\}$	$\{(1,0)\}$	$\{(0,2)\}$	$\{(0,0)\}$	$\{\}$
$\mathcal{G}^{\text{recv}}(n \cdot \texttt{ppn} + p)$	$\{(0,2)\}$	$\{(0,1)\}$	$\{\}$	$\{(0,0)\}$	$\{(1,1)\}$	$\{(1,0)\}$

Table X: Inter-node communication requirements of each process (p, n) for the example given in Figure 4

.

let $\mathcal{I}_{\text{send}}(n \cdot \texttt{ppn} + p, m \cdot \texttt{ppn} + q)$, which defines the global data indices sent from process (p, n) to (q, m):

$$\mathcal{I}(n \cdot \texttt{ppn} + p, m \cdot \texttt{ppn} + q) = \{\mathcal{E}^n_m \,|\, m \text{ maps to } (p, n)\} \tag{24}$$

The global vector indices to which each process (p, n) sends and receives for the example in Figure 4 are displayed in Table XI.

4.2. Local Communication

The function $\mathcal{G}_{\text{send}}(n \cdot \texttt{ppn} + p)$ for $p = 0, \ldots, \texttt{ppn} - 1$, describes evenly distributed inter-node communication requirements for all processes local to node n. However, many of the vector indices to be sent to off-node process $(q, m) \in \mathcal{D}(n \cdot \texttt{ppn} + p, m \cdot \texttt{ppn} + q)$, are not stored on process (p, n). For instance, in Example XI, process $(0, 1)$ sends global vector indices 0 and 1. However, only row 1 is stored on process $(0, 1)$, requiring row 0 to be communicated before inter-node messages are sent.

Similarly, many of the indices that a process (q, m) receives from (p, n) are redistributed to various processes on node n. Example XI requires process $(1, 2)$ to receive

(q,m) \ (p,n)	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
(0, 0)	{}	{}	{}	{}	$\{4,5\}$	{}
(1, 0)	{}	{}	$\{3\}$	{}	{}	{}
(0, 1)	{}	{}	{}	{}	{}	{}
(1, 1)	$\{0\}$	{}	{}	{}	{}	{}
(0, 2)	{}	{}	{}	$\{2\}$	{}	{}
(1, 2)	{}	$\{0,1\}$	{}	{}	{}	{}

Table XI: Inter-node communication pattern for each set of processes (p,n) and (q,m). Each column (p,n) contains the values sent from (p,n) to (q,m), while each row (q,m) contains the values process (q,m) receives from (p,n).

vector data according to indices 0 and 1. Process $(0,2)$ uses both of these vector values, yielding a requirement for redistribution of data received from inter-node communication. Therefore, local communication requirements must be defined.

Each TAPSpMV consists of multiple steps of intra-node communication. Let a function $\mathcal{L}_{\text{send}}(n \cdot \texttt{ppn} + p, \texttt{locality})$ define all processes, local to node n, to which process (p,n) sends messages, where `locality` is a tuple describing the locality of both the original location of the data as well as its final destination. Furthermore, let a function $\mathcal{L}^{\text{recv}}(m \cdot \texttt{ppn} + q, \texttt{locality})$ return the processes (s,m) from which process (q,m) receives. For both functions, the locality of each position is described as either `on_node`, meaning a process local to node n, or `off_node`, meaning a process local to node $m \neq n$.

There are three possible combinations for `locality`: 1) the data is initialized `on_node` with a final destination `off_node`; 2) the original data is `off_node` while the final destination is `on_node`; or 3) both the original data and the final location are `on_node`. These three types of intra-node communication are described in more detail in the remainder of Section 4.2.1.

For each process $(s,n) \in \mathcal{L}(n \cdot \texttt{ppn} + p, \texttt{locality})$, $\mathcal{J}(n \cdot \texttt{ppn} + p, n \cdot \texttt{ppn} + s, \texttt{locality})$ define the global vector indices to be sent from process (p,n) to (s,n) through intra-node communication. This notation is used in following sections.

4.2.1. Local redistribution of initial data. During inter-node communication, a process (p,n) sends all vector values corresponding to the global indices in $\mathcal{J}(n \cdot \texttt{ppn} + p, m \cdot \texttt{ppn} + q)$ to each process $(q,m) \in \mathcal{G}_{\text{send}}(n \cdot \texttt{ppn} + p)$. The indices in $\mathcal{J}(n \cdot \texttt{ppn} + p, m \cdot \texttt{ppn} + q)$ originate on node n, but not necessarily process (p,n). Therefore, the initial vector values but be redistributed among all processes local to node n.

Let $\mathcal{L}_{\text{send}}(n \cdot \texttt{ppn} + p, (\texttt{on_node}, \texttt{off_node}))$ represent all processes, local to node n, to which (p,n) sends initial vector values. Furthermore, let $\mathcal{L}^{\text{recv}}(m \cdot \texttt{ppn} + q, (\texttt{on_node}, \texttt{off_node}))$ be all processes, local to node n, from which (q,m) must receive initial values before intra-node communication can happen. These functions are defined as

$$\mathcal{L}_{\text{send}}(n \cdot \texttt{ppn} + p, (\texttt{on_node}, \texttt{off_node})) = \{(s,n) \,|\, \exists\, j \in \mathcal{R}(n \cdot \texttt{ppn} + p) \ni j \in \mathcal{J}((s,n), m \cdot \texttt{ppn} + q)\} \quad (25)$$

and

$$\mathcal{L}^{\text{recv}}(n \cdot \texttt{ppn} + s, (\texttt{on_node}, \texttt{off_node})) = \{(p,n) \,|\, \exists\, j \in \mathcal{R}(n \cdot \texttt{ppn} + p) \ni j \in \mathcal{J}((s,n), m \cdot \texttt{ppn} + q)\}\,. \quad (26)$$

The local processes to which each (p,n) sends initial data for the example in Figure 4 are displayed in Table XII. Furthermore, the data global vector indices that must be

	(p,n) (0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
$\mathcal{L}_{\text{send}}$	$\{(1,0)\}$	$\{\}$	$\{(1,1)\}$	$\{(0,1)\}$	$\{\}$	$\{(0,2)\}$
$\mathcal{L}^{\text{recv}}$	$\{\}$	$\{(0,0)\}$	$\{(1,1)\}$	$\{(0,1)\}$	$\{(1,2)\}$	$\{\}$

Table XII: Initial intra-node communication requirements for each process (p,n) for the example in Figure 4. The rows of the table describe $\mathcal{L}_{\text{send}}(n\cdot\texttt{ppn}+p,(\texttt{on_node},\texttt{off_node}))$ and $\mathcal{L}^{\text{recv}}(n\cdot\texttt{ppn}+p,(\texttt{on_node},\texttt{off_node}))$.

sent from process (p,n) to each $(s,n)\in\mathcal{L}_{\text{send}}(n\cdot\texttt{ppn}+p,(\texttt{on_node},\texttt{off_node}))$, and the indices of process (s,n) receives from each $(p,n)\in\mathcal{L}^{\text{recv}}(n\cdot\texttt{ppn}+s,(\texttt{on_node},\texttt{off_node}))$ are defined as

$$\mathcal{J}(n\cdot\texttt{ppn}+p,n\cdot\texttt{ppn}+s,(\texttt{on_node},\texttt{off_node}))= \{i\,|\,i\in\mathcal{R}(n\cdot\texttt{ppn}+p),\forall\, i\in\mathcal{G}_{\text{send}}(n\cdot\texttt{ppn}+s)\}\,. \quad (27)$$

The global vector indices that each (p,n) must send to other processes on node n for the example in Figure 4 are displayed in Figure XIII.

	(q,n) \ (p,n)	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
	(0, 0)	$\{\}$	$\{\}$	—	—	—	—
	(1, 0)	$\{0\}$	$\{\}$	—	—	—	—
	(0, 1)	—	—	$\{\}$	$\{3\}$	—	—
	(1, 1)	—	—	$\{2\}$	$\{\}$	—	—
	(0, 2)	—	—	—	—	$\{\}$	$\{5\}$
	(1, 2)	—	—	—	—	$\{\}$	$\{\}$

Table XIII: Global vector indices of initial data that is communicated between processes local to each node n for the example in Figure 4. Each column contains the indices of values sent from (p,n) to (q,n), while each row shows the indices of values that (q,m) receives from (p,n). Note: dashes (—) throughout the table represent processes on separate nodes, which do not communicate during intra-node communication.

4.2.2. Local redistribution of received off-node data. For each receive in $\mathcal{G}^{\text{recv}}(m\cdot\texttt{ppn}+q)$, process (q,m) receives all values needed by any process on node n. These values are then distributed across the processes local to n. Let $\mathcal{L}_{\text{send}}(m\cdot\texttt{ppn}+q,(\texttt{off_node},\texttt{on_node}))$ define all processes local to node m to which process (q,m) sends vector values that have been received through inter-node communication. Furthermore, $\mathcal{L}^{\text{recv}}(m\cdot\texttt{ppn}+s,(\texttt{off_node},\texttt{on_node}))$ define all processes from which (s,m) receives vector values that originated off-node. These functions are defined as

$$\mathcal{L}_{\text{send}}(m\cdot\texttt{ppn}+q,(\texttt{off_node},\texttt{on_node}))= \{(s,m)\,|\,\exists\, A_{ij}\neq 0 \text{ with } i\in\mathcal{R}(m\cdot\texttt{ppn}+s),j\in\mathcal{J}(n\cdot\texttt{ppn}+p,m\cdot\texttt{ppn}+q)\} \quad (28)$$

and

$$\mathcal{L}^{\text{recv}}(m\cdot\texttt{ppn}+s,(\texttt{off_node},\texttt{on_node}))= \{(q,m)\,|\,\exists\, A_{ij}\neq 0 \text{ with } i\in\mathcal{R}(m\cdot\texttt{ppn}+s),j\in\mathcal{J}(n\cdot\texttt{ppn}+p,m\cdot\texttt{ppn}+q)\}\,. \quad (29)$$

These are highlighted, for the example in Figure 4, in Table XIV. Furthermore, the data global vector indices that must be sent from process (q,m) to each $(s,m)\in$

	(p,n) (0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
$\mathcal{L}_{\text{send}}$	{}	$\{(0,0)\}$	{}	{}	{}	$\{(0,2)\}$
$\mathcal{L}^{\text{recv}}$	$\{(1,0)\}$	{}	{}	{}	$\{(1,2)\}$	{}

Table XIV: Intra-node communication requirements containing processes to which each (p,n) sends received inter-node data, according to the example in Figure 4. The rows of the table describe $\mathcal{L}_{\text{send}}(n \cdot \texttt{ppn} + p, (\texttt{off_node}, \texttt{on_node}))$ and $\mathcal{L}^{\text{recv}}(n \cdot \texttt{ppn} + p, (\texttt{off_node}, \texttt{on_node}))$.

$\mathcal{L}_{\text{send}}(m \cdot \texttt{ppn} + q, (\texttt{off_node}, \texttt{on_node}))$, as well as the indices that some process (s,m) receives from each $(q,m) \in \mathcal{L}^{\text{recv}}(m \cdot \texttt{ppn} + s, (\texttt{off_node}, \texttt{on_node}))$ are defined as

$$\mathcal{J}(m \cdot \texttt{ppn} + q, m \cdot \texttt{ppn} + s, (\texttt{off_node}, \texttt{on_node})) = \{j \in \mathcal{I}(n \cdot \texttt{ppn} + p, m \cdot \texttt{ppn} + q) \,|\, A_{ij} \neq 0 \text{ with } i \in \mathcal{R}(m \cdot \texttt{ppn} + s)\}\,. \quad (30)$$

The global vector indices, received from the inter-node communication step, which (p,n) must send to each local process (q,n) in the example in Figure 4 are displayed in Table XV.

(q,n) \ (p,n)	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
(0, 0)	{}	{3}	—	—	—	—
(1, 0)	{}	{}	—	—	—	—
(0, 1)	—	—	{}	{}	—	—
(1, 1)	—	—	{}	{}	—	—
(0, 2)	—	—	—	—	{}	{1}
(1, 2)	—	—	—	—	{}	{}

Table XV: Global vector indices of received inter-node data that must be communicated between processes local to each node n for the example in Figure 4. Each column contains the indices of values sent from (p,n) to (q,n), while each row shows the indices of values that (q,m) receives from (p,n). Note: dashes (—) throughout the table represent processes on separate nodes, which cannot communicate during intra-node communication.

4.2.3. Fully Local Communication. A subset of the values needed by a process (p,n) are stored on local process (s,n). One advantage is that these values bypass the three-step communication, and are communicated directly. Let $\mathcal{L}_{\text{send}}(n \cdot \texttt{ppn} + p, (\texttt{on_node}, \texttt{on_node}))$ define all processes local to node n to which (p,n) sends vector data. Furthermore, let $\mathcal{L}^{\text{recv}}(n \cdot \texttt{ppn} + s, (\texttt{on_node}, \texttt{on_node}))$ define all processes local to node n from which (s,n) receives vector data. These functions are defined as

$$\mathcal{L}_{\text{send}}(n \cdot \texttt{ppn} + p, (\texttt{on_node}, \texttt{on_node})) = \{(s,n) \,|\, \exists\, A_{ij} \neq 0 \text{ with } i \in \mathcal{R}(n \cdot \texttt{ppn} + p), j \in \mathcal{R}(n \cdot \texttt{ppn} + s)\} \quad (31)$$

and

$$\mathcal{L}^{\text{recv}}(n \cdot \texttt{ppn} + s, (\texttt{on_node}, \texttt{on_node})) = \{(p,n) \,|\, \exists\, A_{ij} \neq 0 \text{ with } i \in \mathcal{R}(n \cdot \texttt{ppn} + p), j \in \mathcal{R}(n \cdot \texttt{ppn} + s)\}\,. \quad (32)$$

The processes, local to node n, to which (p,n) must send initial vector data for the example in Figure 4 are displayed in Table XVI. Furthermore, the data global vec-

	(p,n)					
	(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
$\mathcal{L}_{\text{send}}$	$\{\}$	$\{(0,0)\}$	$\{\}$	$\{(0,1)\}$	$\{\}$	$\{\}$
$\mathcal{L}^{\text{recv}}$	$\{(1,0)\}$	$\{\}$	$\{(1,1)\}$	$\{\}$	$\{\}$	$\{\}$

Table XVI: Intra-node communication requirements containing processes to which each process (p,n) must send vector data, according to the example in Figure 4. The rows of the table describe $\mathcal{L}_{\text{send}}(n \cdot \texttt{ppn} + p, (\texttt{on_node}, \texttt{on_node}))$ and $\mathcal{L}^{\text{recv}}(n \cdot \texttt{ppn} + p, (\texttt{on_node}, \texttt{on_node}))$.

tor indices that must be sent from process (p,n) to each $(s,n) \in \mathcal{L}_{\text{send}}(n \cdot \texttt{ppn} + p, (\texttt{on_node}, \texttt{on_node}))$, as well as the indices that some process (s,n) receives from each $(p,n) \in \mathcal{L}^{\text{recv}}(n \cdot \texttt{ppn} + s, (\texttt{on_node}, \texttt{on_node}))$ are defined as follows.

$$\mathcal{J}(n \cdot \texttt{ppn} + p, n \cdot \texttt{ppn} + s, (\texttt{on_node}, \texttt{on_node})) = \{j \,|\, \exists\, A_{ij} \neq 0 \text{ with } i \in \mathcal{R}(n \cdot \texttt{ppn} + p), j \in \mathcal{R}(n \cdot \texttt{ppn} + s)\} \quad (33)$$

The global vector indices which (p,n) must send to each local process (q,n) in the example in Figure 4 are displayed in Table XVII.

		(p,n)					
		(0, 0)	(1, 0)	(0, 1)	(1, 1)	(0, 2)	(1, 2)
(q,n)	(0, 0)	$\{\}$	$\{1\}$	—	—	—	—
	(1, 0)	$\{\}$	$\{\}$	—	—	—	—
	(0, 1)	—	—	$\{\}$	$\{3\}$	—	—
	(1, 1)	—	—	$\{\}$	$\{\}$	—	—
	(0, 2)	—	—	—	—	$\{\}$	$\{\}$
	(1, 2)	—	—	—	—	$\{\}$	$\{\}$

Table XVII: Global vector indices that must be communicated between processes local to each node n for the example in Figure 4. Each column contains the indices of values sent from (p,n) to (q,n), while each row shows the indices of values that (q,m) receives from (p,n). Note: dashes (—) throughout the table represent processes on separate nodes, which cannot communicate during intra-node communication.

4.3. Alternative SpMV Algorithm

5. RESULTS

In this section, the parallel performance and scalability of the TAPSpMV in comparison to the standard SpMV is presented. Scaling tests are considered for a structured 2D rotated anisotropic diffusion and for unstructed linear elasicity. In addition, the matrix-vector multiplication in an algebraic multigrid (AMG) hierarchy is tested in order to expose a variety of communication patterns. Lastly, scaling tests on the largest 15 matrices from the Florida sparse matrix collection are presented. All tests are performed on the Blue Waters parallel computer at University of Illinois at Urbana-Champaign. AMG hierarchies consist of successively coarser, but denser, levels. Therefore, while a standard SpMV performed on the original matrix often requires communication of a small number of large messages, coarse levels require a large number of small messages to be injected into the network. Figure 6 shows that both the number and size of inter-node messages required on each level of the linear elasticity hierarchy are reduced through use of the TAPSpMV. There is a large reduction in communication requirements for coarse levels of the hierarchy, which includes a high number of small

ALGORITHM 2: `local_communication`

Input: (p, n) : tuple describing local rank and node of process
$v|_{\mathcal{R}(n\cdot \text{ppn}+p)}$: rows of input vector v local to process (p, n)
`locality`: locality of input and output data

Output: ℓ_{recv}: values that rank (p, n) receives from other processes

for $(s, n) \in \mathcal{L}_{send}(n \cdot \text{ppn} + p, \texttt{locality})$ **do**
  for $i \in \mathcal{J}(n\text{ppn} + p, n\text{ppn} + s, \texttt{locality})$ **do**
    $\ell_{\text{send}} \leftarrow v|_{\mathcal{R}(n\cdot \text{ppn}+p)_i}$
  end
  `MPI_Isend`$(\ell_{\text{send}}, \ldots, n \cdot \text{ppn} + s, \ldots)$ `// Initialize sends`
end
$\ell_{\text{recv}} \leftarrow \emptyset$
for $(s, n) \in \mathcal{L}^{recv}(n \cdot \text{ppn} + p, \texttt{locality})$ **do**
  `MPI_Irecv`$(\ell_{\text{recv}}, \ldots, n \cdot \text{ppn} + s, \ldots)$ `// Initialize receives`
end
`MPI_Waitall` `// Complete sends and receives`

ALGORITHM 3: `new_spmv`

Input: (p, n): tuple describing local rank and node of process
$A|R$: rows of matrix A local to process (p, n)
$v|R$: rows of input vector v local to process (p, n)

Output: $w|\mathcal{R}$: rows of output vector $w \leftarrow Av$, local to process (p, n)

$A_{\text{on_process}} = \texttt{on_process}(A|\mathcal{R})$
$A_{\text{on_node}} = \texttt{on_node}(A|\mathcal{R})$
$A_{\text{off_node}} = \texttt{off_node}(A|\mathcal{R})$
$b_{\ell\to\ell} \leftarrow \texttt{local_communication}(n \cdot \text{ppn} + p, v|\mathcal{R}, (\texttt{on_node} \to \texttt{off_node}))$
$b_{\ell\to n\ell} \leftarrow \texttt{local_communication}(n \cdot \text{ppn} + p, v|\mathcal{R}, (\texttt{on_node} \to \texttt{on_node}))$
for $(q, m) \in \mathcal{G}_{send}(n\text{ppn} + p)$ **do**
  for $i \in \mathcal{I}(n\text{ppn} + p, m\text{ppn} + q)$ **do**
    $g_{\text{send}} \leftarrow b^i_{\ell\to n\ell}$
  end
  `MPI_Isend`$(g_{\text{send}}, \ldots, m\text{ppn} + q, \ldots)$ `// Initialize sends`
end
$g_{\text{recv}} \leftarrow \emptyset$
for $(q, m) \in \mathcal{G}^{recv}(n\text{ppn} + p)$ **do**
  `MPI_Irecv`$(g_{\text{recv}}, \ldots, m\text{ppn} + q, \ldots)$ `// Initialize receives`
end
`sequential_spmv`$(A_{\text{on_process}}, v|\mathcal{R})$ `// Sequential SpMV for local values`
`sequential_spmv`$(A_{\text{on_node}}, b_{\ell\to\ell})$ `// Sequential SpMV for on-node values`
`MPI_Waitall` `// Complete sends and receives`
$b_{n\ell\to\ell} \leftarrow \texttt{local_communication}(n \cdot \text{ppn} + p, v|\mathcal{R}, (\texttt{off_node} \to \texttt{on_node}))$
`sequential_spmv`$(A_{\text{off_node}}, b_{n\ell\to\ell})$ `// Sequential SpMV for off-node values`

messages. However, as the TAPSpMV requires redistribution of data among processes local to each node, the intra-node communication requirements increase greatly for the TAPSpMV, as shown in Figure 7.

While there is an increase in intra-node communication requirements, the reduction in more expensive inter-node messages results in a significant reduction in total time

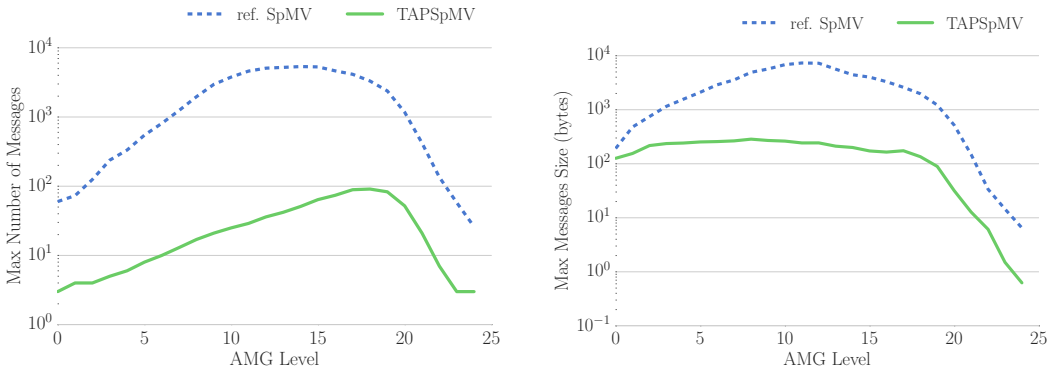

Fig. 6: The maximum number (left) and size (right) of **inter-node** messages communicated by a single process during a standard SpMV and TAPSpMV on each level of the **linear elasticiy AMG hierarchy**.

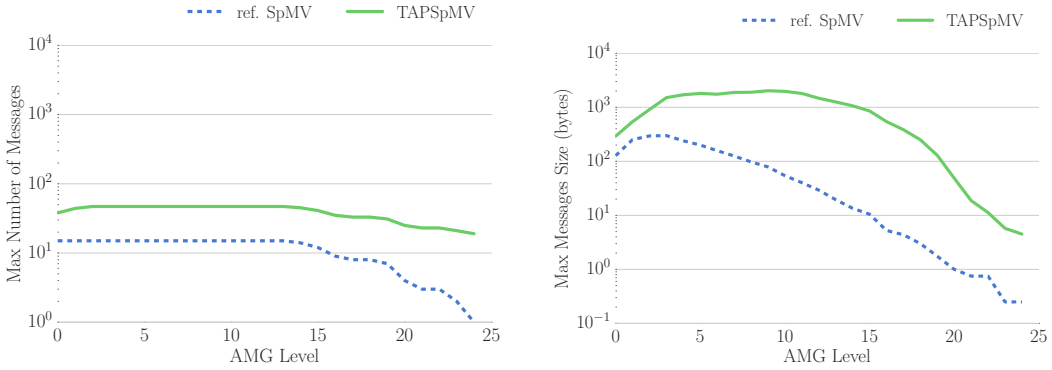

Fig. 7: The maximum number (left) and size (right) of **intra-node** messages communicated by a single process during a standard SpMV and TAPSpMV on each level of the **linear elasticiy AMG hierarchy**.

for the TAPSpMV algorithm, particularly on coarser levels near the middle of each AMG hierarchy.

6. CONCLUSION AND FUTURE WORK

This paper introduces a method to reduce communication that is injected into the network during a sparse matrix-vector multiply by reorganizing messages on each node. This results in a reduction of the inter-node communication, replaced by less-costly intra-node communication, which reduces both the number and size of messages that are injected into the network.

REFERENCES

C. Chevalier and F. Pellegrini. 2008. PT-Scotch: A tool for efficient parallel graph ordering. *Parallel Comput.* 34, 68 (2008), 318 – 331. DOI:http://dx.doi.org/10.1016/j.parco.2007.12.001 Parallel Matrix Algorithms and Applications.

William Gropp, Luke N. Olson, and Philipp Samfass. 2016. Modeling MPI Communication Performance on SMP Nodes: Is It Time to Retire the Ping Pong Test. In *Proceedings of the*

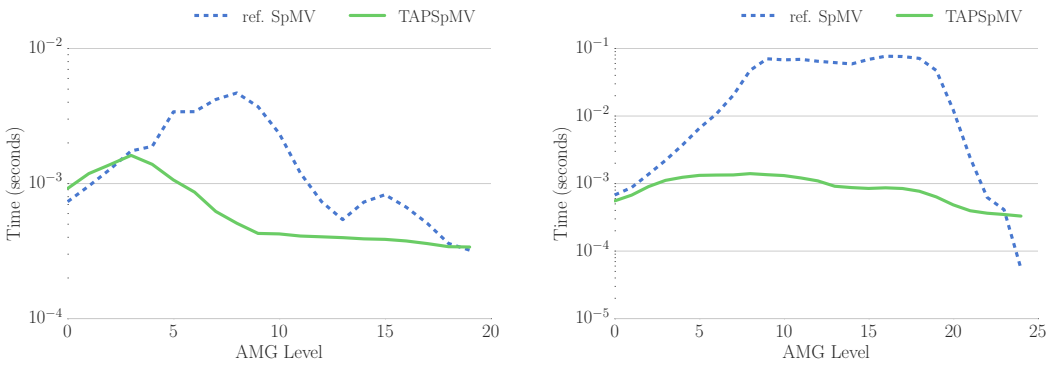

Fig. 8: The time required to perform the various SpMVs on each level of the rotated anisotropic (left) and linear elasticity (right) AMG hierarchies.

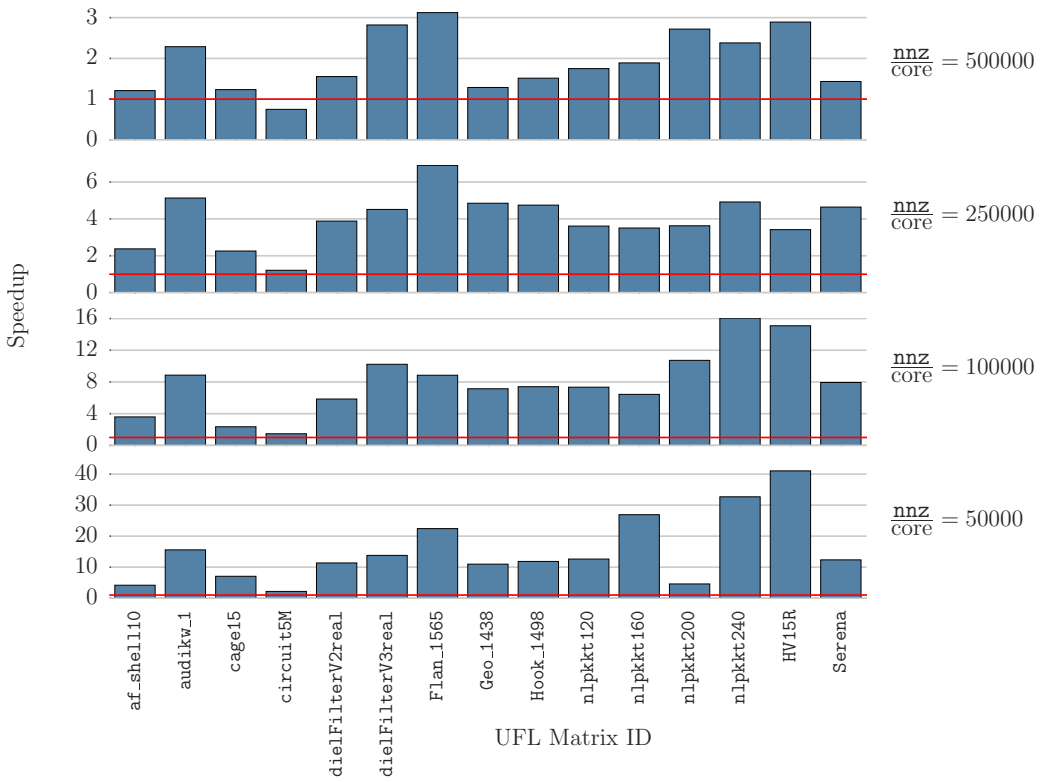

Fig. 9: The time required to perform the standard and TAPSpMVs on a subset of the 15 largest real matrices from the Florida sparse matrix collection at various scales, were nz is the average number of nonzeros per core, partitioned so that each row r is stored on process $p = r \mod n_p$, where n_p is the number of processes.

23rd European MPI Users' Group Meeting (EuroMPI 2016). ACM, New York, NY, USA, 41–50. DOI:http://dx.doi.org/10.1145/2966884.2966919

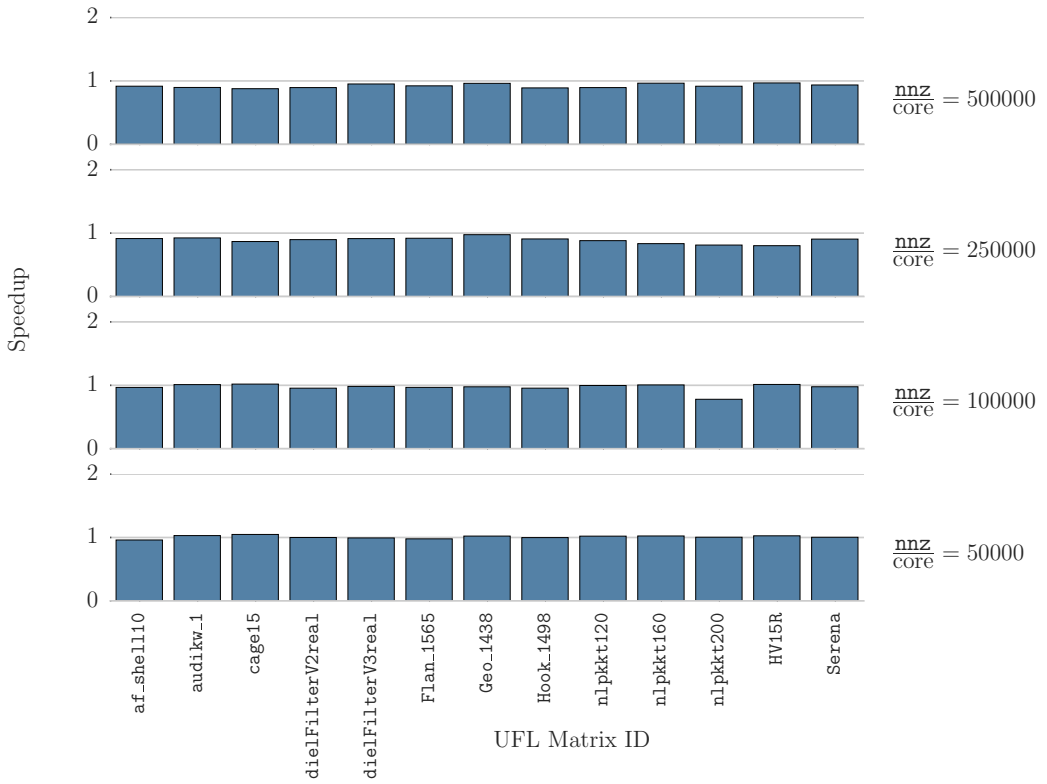

Fig. 10: The time required to perform the standard and TAPSpMVs on a subset of the 15 largest real matrices from the Florida sparse matrix collection at various scales, where nz is the average number of nonzeros per core, partitioned with PT Scotch. **Note: the matrices with ID 1905 and 2276 are not included in this subset, due to partitioning constraints.**

George Karypis and Vipin Kumar. 1998. A Parallel Algorithm for Multilevel Graph Partitioning and Sparse Matrix Ordering. *J. Parallel Distrib. Comput.* 48, 1 (1998), 71–95. DOI:http://dx.doi.org/10.1006/jpdc.1997.1403